



RIDGE SECURITY

RidgeOne

| *AI offensive security in one box*

DATASHEET

RIDGE SECURITY, INC.

RidgeOne brings autonomous penetration testing and agentic red teaming into a single appliance. It combines RidgeGen, Ridge Security's agentic security validation platform, with an open-weight AI model running on NVIDIA DGX Spark — advanced AI security testing with no AI infrastructure to build, no AI team to hire, and no per-token costs.

01 EXECUTIVE SUMMARY

Attackers now use AI to find and exploit weaknesses faster than most security teams can test for them. Defending against that takes the same class of technology on your side — and for many organizations, getting it has been the hard part. Cloud-based AI security tools need reliable frontier-model access, a budget that absorbs per-query charges, and a team to run the arrangement. Building the capability in-house needs AI infrastructure and expertise most security teams should not have to acquire.

RidgeOne removes that barrier. It is a single appliance that arrives ready to run: RidgeGen and an open-weight AI model, pre-integrated on NVIDIA DGX Spark. You plug it in, define what to test, and it goes to work. No AI infrastructure to build, no model to source or tune, and no per-token charge from an AI provider.

What RidgeOne gives you

Advanced testing, ready to run

Simple deployment

Predictable cost

Deep testing and remediation

Portable

Your data stays with you

Why it matters

Agentic AI penetration testing without an AI team or AI infrastructure.

Import the appliance, connect it, start testing — five steps.

One compact appliance; testing volume generates no AI-provider bill.

Findings are verified by exploitation, linked into attack paths, and come with fix guidance.

Small and light enough to move between sites and network segments.

Testing data and AI reasoning stay on the appliance; no outbound connection is needed to run a test.

RidgeOne does not replace RidgeGen software deployments. Organizations that need large-scale or parallel testing, or deployment into a private cloud or data center they already run, can continue to run RidgeGen as software and scale it with their infrastructure.

02 WHO IT IS FOR, AND WHERE IT FITS

RidgeOne suits organizations that want modern AI security testing without the overhead that usually comes with it — defined by workload and deployment need rather than company size. It is the right fit when frontier-model access is limited or inconvenient, your team has no AI specialists, your testing

workload is steady rather than enormous, predictable cost matters, or you would rather testing data never left your premises. RidgeOne and RidgeGen software run the same platform; what differs is how you buy it, deploy it, and the scale it is sized for.

Consideration	RidgeOne	RidgeGen software
Best suited to	Mid-market, steady workloads	Larger, parallel or custom needs
Deployment	Pre-integrated appliance	Your own managed infrastructure
Infrastructure	NVIDIA DGX Spark included	You provide supported infrastructure
AI model	Open-weight, packaged and ready	Aligned to your environment
Scaling	Four concurrent assessments	Scales with your infrastructure
Why customers choose it	Simple acquisition, predictable cost	Deployment flexibility and scale

03 WHAT RIDGEONE TESTS

RidgeOne performs autonomous penetration testing and agentic red teaming across the environment — not just web applications. It works the way an experienced tester does: it explores, forms a view of what is likely weak,

tries it, and uses what it learns to decide what to try next. Findings are confirmed by exploiting them, so what you receive has been demonstrated rather than guessed at, and mapped to OWASP Top 10 and MITRE ATT&CK.

- **IT infrastructure:** Hosts, services, configuration weaknesses and exposure
- **Networks:** Internal segments, reachable services and lateral-movement paths
- **Web applications:** Business logic, access control, injection and authentication flaws
- **APIs:** Authorization, object-level access control and data exposure
- **Active Directory:** Identity paths, privilege escalation and credential weaknesses
- **Hosts, IoT and OT:** OS- and service-level weaknesses, and connected/operational assets in scope

04 INSIDE THE BOX

Hardware

Component	Specification
System	NVIDIA DGX Spark (GB10 Grace Blackwell)
Memory	128 GB unified, shared between the AI model and RidgeGen
Processors	20-core Arm CPU with integrated Blackwell GPU
Storage	4 TB NVMe, self-encrypting
Networking	10 GbE plus a high-speed network adapter
Power / size	240 W, standard office or rack power — 150 × 150 × 50 mm, 1.2 kg

Software and AI models

RidgeOne runs Ridge Security’s RidgeGen platform with an open-weight AI model on the appliance’s GPU — reachable only from inside the box. RidgeGen is model-agnostic, and RidgeOne now ships the NVIDIA Nemotron 3 family of open models, a natural fit for the NVIDIA DGX Spark. You choose the model that matches your workload; the platform, your scope and your safety controls work the same way regardless of which one runs.

AI model	Scale	Best for
NVIDIA Nemotron 3 Nano	30B	Efficient, fast turnaround — the default for steady testing throughput
NVIDIA Nemotron 3 Super	120B (MoE)	Deeper reasoning for complex, multi-step targets
GPT-OSS	Open-weight	Also supported — RidgeGen is model-agnostic


Open-weight by design

No per-seat model licensing, no usage caps, and a model your own team can inspect. The appliance runs Ubuntu 24.04 LTS Server hardened to CIS Benchmark Level 1, is administered through a web interface plus a restricted management console, and ships as a signed image with published checksums. The exact model build for each release is stated in the release documentation, and the model can be updated as open-weight models improve without changing how the platform or your safety controls work.

05 HOW IT WORKS

RidgeOne runs two pieces of software on one appliance. The AI model does the thinking; RidgeGen does the testing. The model runs on the GPU — it decides what to test next and interprets what comes back. RidgeGen runs on the CPU — it executes the tests, applies your safety rules, and writes the report. The two take

turns inside the box: the model decides, the platform acts, the model reads the result and decides again. No request leaves the appliance for the AI to do its work, and the platform — not the model — enforces the boundaries you set, keeping reasoning and execution control separate.



06 DEPLOYMENT IN FIVE STEPS

Step	What you do	Result
1) Verify	Download the signed image and check the published checksum	You know the image is genuine
2) Import	Deploy into your VMware environment and connect it to the network under test	The appliance is on your network
3) First boot	It provisions its own credentials and certificates	Each appliance has a unique identity
4) Configure	Sign in, set an admin account, define scope and testing posture	Boundaries are set before anything runs
5) Test	Start the assessment	Testing runs on its own and produces a report

There are no model downloads, GPU drivers or AI servers to set up — that work is done before the appliance reaches you.

07 GETTING STARTED

Ridge Security will scope a proof of value against systems you select, in your own environment, so you can judge RidgeOne on your applications rather than on ours. A typical evaluation covers a defined scope in one network segment, runs an assessment end to end, and produces the same evidence-backed report your team would receive in production. We will also walk through sizing, deployment, and how RidgeOne fits the rest of your testing program.

IN SHORT

RidgeOne combines RidgeGen, Ridge Security's agentic security validation platform, with an open-weight AI model — the NVIDIA Nemotron 3 family (Nano and Super) or GPT-OSS — running on NVIDIA DGX Spark. It performs autonomous penetration testing and agentic red teaming across IT infrastructure, networks, web, APIs, Active Directory and hosts; and keeps your testing data on hardware you control.

Ridge Security delivers autonomous cybersecurity validation that helps organizations manage risk and improve resilience, and is a member of the NVIDIA Inception program. To arrange an evaluation, contact your Ridge Security representative or visit ridgesecurity.ai.

